Finite time distributed averaging over ring networks

Alessandro Falsone, Kostas Margellos, Simone Garatti, Maria Prandini

arXiv:1608.08358v1 [math.OC] 30 Aug 2016

***Abstract*—We consider a multi-agent system where each agent has its own estimate of a given quantity and the goal is to reach consensus on the average. To this purpose, we propose a distributed consensus algorithm that guarantees convergence to the average in a finite number of iterations. The algorithm is tailored to ring networks with bidirectional pairwise communications. If the number of agents m is even, say $m = 2n$, then, the number of iterations needed is equal to n, which in this case is the diameter of the network, whereas the number of iterations grows to $3n$ if the number of agents is odd and equal to $m = 2n + 1$.**

***Index Terms*—Consensus, gossip algorithms, distributed averaging, networks.**

I. Introduction

A typical problem encountered in a multi-agent system is that all agents are aiming at some common goal but they can communicate only with their neighbors to this purpose. Achieving a common goal often translates into reaching an agreement on the value taken by some quantity (the agreement variable) via some consensus algorithm. Basic consensus algorithms date back to [1]. Only in the last two decades, they have attracted the attention of both the computer science [2] and control engineering [3] communities.

There are different forms of consensus. Here, we are concerned with distributed averaging where each agent has its own estimate of a certain quantity and the goal is to reach consensus on the average of the values stored by all agents. The two most common approaches in the literature to the distributed averaging problem are linear iterations and gossip algorithms. The former is an iterative scheme where each agent updates its estimate of the average by taking a linear combination of its current estimate and those received by its neighboring agents, [4]. Gossip algorithms are similar, but they require that only pairwise communication occurs, [5], [6]. In both approaches, the interactions between agents can be modeled as a time-varying weighted graph, with vertices and edges respectively representing the agents and the communication links, and weights assigned to edges being the coefficients of the linear combinations. Under mild assumptions on the coefficients and on the structure of the graph across iterations, both linear iterations schemes and gossip algorithms have been proven to asymptotically converge to the average, [4]. Since reaching consensus asymptotically can be limiting in practice, there is a fair amount of literature on how to analyze and optimize the convergence rate of both algorithms, see [4], [7] and references therein.

As pointed out in [2], [4], [6], distributed average computation can be actually solved in a finite number of iterations by making each agent keep collecting all the values received and passing them to its neighbors. After a number of iterations equal to the diameter of the graph (i.e., the maximum distance between any two vertices), every agent knows all the values and can compute their average. This algorithm is apparently very simple but it actually presents two main issues: i) the number of values that each agent needs to store (and thus the memory usage) grows linearly with the number of agents, and ii) an unnecessary amount of information is exchanged which might overload the communication channels. If the quantities to be averaged were vectors instead of scalars, the previous two issues would become even more critical. Much effort has then be devoted in the literature to design simple algorithms (like linear iterations) that are able to solve the distributed averaging problem in finite time, see [8], [9], [10], [11], [12], [13], [14], [15] just to name a few. Most of these works take a global perspective, in that they assume to know the topology of the (undirected) graph, and then orchestrate the weights so as to achieve finite time convergence. In [8] the topology is assumed to be time-invariant, and the weights are kept constant across iterations. The solution to the finite time averaging problem is then given in terms of the minimal polynomial of some matrix which gathers the weights and matches the graph topology. Convergence is achieved in $D+1$ steps, where D is the degree of the minimal polynomial, provided that each agent stores all the $D+1$ previous values, which, however, might be impractical for large networks. A further result in [8] is that the degree of the minimal polynomial can be computed in a distributed way in a number of iterations that is equal to the number of agents. In [11] the network is also assumed to be time-invariant, but the weights (and thus the weight matrix) change across iterations. An analytic solution based on the joint diagonalization of the weight matrices is provided, and it is shown that convergence can be achieved in a number of steps equal to the number of distinct eigenvalues of the Laplacian matrix of the graph.

In [15] the authors consider a fixed network topology and develop an algorithm based on a two-stage max-consensus which achieves finite time convergence in $d(2m+1)$ iterations, where d is the diameter of the graph and m is the number of agents. In [9], [13], [14], the authors focus only on trees, i.e., graphs without loops, and show how to reach finite time convergence over these topologies. Note that, even though it is always possible to construct a spanning tree (i.e., a tree which reaches all vertices) from a connected graph by dropping some

Corresponding author: Alessandro Falsone.

Alessandro Falsone, Simone Garatti, and Maria Prandini are with Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Milano, Italy (phone: +39 02 2399 {4028,3650,3441} e-mail: {alessandro.falsone, simone.garatti, maria.prandini}@polimi.it).

Kostas Margellos is with Department of Engineering Science, University of Oxford, Oxford, United Kingdom (phone: +44 1865 283912; e-mail: kostas.margellos@eng.ox.ac.uk).

This work is partially supported by the European Commission under the project UnCoVerCPS with grant number 643921.

links, the diameter of the resulting tree might be greater than the diameter of the original graph, thus requiring more iterations to reach convergence. In [10] the finite time averaging problem is rephrased as a nonlinear optimization program that can be solved numerically, where the weights in the graph are optimization variables. Finite time average can be achieved in this case with a number of iterations that ranges between d and $2d$, where d is the diameter of the network. However, agents have to solve a nonlinear program in a distributed fashion, which might be difficult from a computational point of view. In [12] the authors prove that for specific type of graphs finite time averaging can be achieved in a number of steps equal to the diameter of the graph. Ring networks are included in this setting, but the communication graph is kept constant and gossip constraints are hence not considered. Finally, in [9], [16], finite time averaging under the gossip constraint has been investigated. In particular, in [16] it is proven that if the weights are constant and equal to $1/2$, then finite time convergence on undirected graphs can be achieved only if the number of vertices equals a power of 2. Finite time averaging can be achieved on a directed graph with weights equal to $1/2$, but in such a case the number of iteration required for a network with $m = 2^n + r$ agents, where $r \in \mathbb{N}$, is $mn + 2r$, which grows more than linearly with m. In [9], it is shown that finite time averaging can be achieved under the gossip constraint by allowing for time-varying coefficients. An upper bound that scales quadratically with the number m of agents is given for the number of iterations required to achieve finite time averaging.

In this work we are concerned with the design of a finite time averaging algorithm specifically tailored to a ring topology, subject to a gossip constraint. By allowing the weights of the communication graph to be time-varying, we are able to prove finite time convergence for a network with an arbitrary number of agents. More specifically, for ring networks with an even number of agents, we propose an algorithm that ensures finite time convergence and attains the lower bound on the number of iterations needed for a general communication graph with no gossip constraint, which is the diameter of the graph. As for ring networks with an odd number of agents, an algorithm based on the case of an even number of vertices is designed. Still convergence in finite time to the average is guaranteed in a number of iterations that scales linearly with the number of agents. In this case, however, the lower bound is not attained. As for memory requirements, each agent needs to store only its estimate of the mean when the total number of agents is even, whereas the memory requirement is doubled in the odd case. In both cases, the required memory is independent of the number of agents.

The rest of the paper is organized as follows. In Section II, we introduce the problem set-up with the ring topology and the gossip constraint. The proposed distributed averaging algorithms for the cases of even and odd number of agents are described in Section III, where finite time convergence and related bounds on the number of iterations are shown. Finally, in Section IV some concluding remarks are drawn.

II. Problem setup

We address the case of a multi-agent system characterized by a ring communication network, where each agent can communicate only with one of each neighbors at a time (gossip constraint) to take some joint decision. The agents are m in total and each one is identified by an integer i taking values in $\{1, 2, \ldots, m\}$. Each agent i has its own estimate $x_i(0)$ of some quantity of interest. The goal is to devise a distributed algorithm through which each agent is able to compute the average

$$\bar{x} = \frac{1}{m} \sum_{i=1}^{m} x_i(0), \tag{1}$$

in a finite number of iterations only by exchanging information with its neighbors under the gossip constraint. Evidently, only the case with $m > 2$ is of interest.

The network communication structure can be represented as an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, \ldots, m\}$ is the set of vertices representing the agents, and $\mathcal{E}$ is the set of edges, defined as the following collection of ordered pairs of vertices:

$$\mathcal{E} = \bigcup_{i=1}^{m-1} \{(i, i+1), (i+1, i)\} \cup \{(1, m), (m, 1)\},$$

according to the ring topology. Since we are dealing with an undirected graph, with a slight abuse of notation we will use (i, j) to denote both edges (i, j) and (j, i). According to this notation, for $m > 2$, the number of edges is even if m is even, odd otherwise.

We next define the communication protocol used by the agents to exchange information over the network. We are here concerned with the design of a synchronous pairwise communication strategy. The agents communicate in rounds, and successive communication rounds are indexed by an integer $k \in \mathbb{N}$, which is also referred to as time step or iteration. At any round k, each agent is allowed to communicate only with one of its neighbors (gossip constraint). Since communication channels are represented by the edges of $\mathcal{G}$, the design of the communication protocol then reduces to specify a proper sequence of edges to activate, where edge (i, j) is said to be *active* at time step k if agents i and j exchange information at time k. Of course, more edges can be activate at the same iteration, as long as they do not have any vertex in common, so as to comply with the pairwise communication constraint (see the multigossip framework in [13]).

To reduce the number of communication rounds we need to parallelize as much as possible the number of simultaneous pairwise communications. This can be interpreted as an edge-coloring problem on $\mathcal{G}$, [13], where edges with the same color represent communication channels which can be active at the same time. For a ring communication network with an even number $m > 2$ of vertices we need just two colors to minimize the number of communication rounds needed for all edges to be activated while satisfying the gossip constraint. In the odd case we need at least three colors. Figure 1 shows the coloring scheme for the even case in a pictorial form, with two groups of edges characterized by two different colors, blue and red,

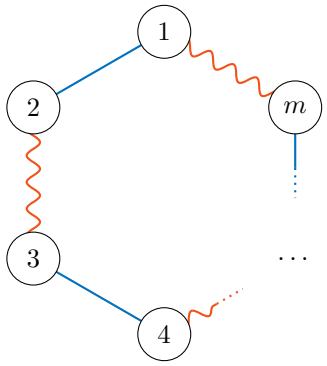

Fig. 1. Communication protocol under the gossip constraint in the case of an even number of agents: edges are grouped in two sets by a color coding. Only edges with the same color are active simultaneously.

that can be activated alternatively, e.g., all the blue straight edges are activate at iterations $k = 1, 3, 5, \ldots$, whereas the red wavy ones at $k = 2, 4, 6 \ldots$.

Finally, we need to specify the weights associated with the edges, which define the linear update of the agents local estimate of the mean based on the information received from its neighbors. To this end, assume that at time k, agents i and j communicate. Agent i performs then the following update for its local estimate of the mean:

$$x_i(k) = (1 - \alpha_k)x_i(k-1) + \alpha_k x_j(k-1), \qquad (2)$$

where α_k is the (time-varying) weight associated with edge (i, j) at time k entering the convex update rule. Agent j performs an analogous update step.

The problem to be addressed is how to appropriately select α_k, $k = 1, 2, \ldots$, so as to guarantee that there exists a finite time T such that $x_i(T) = \bar{x}$ for all $i = 1, \ldots, m$, where $\bar{x}$ is given by (1).

III. Proposed solution

We next provide a solution to the finite time distributed averaging problem over a ring network topology with pairwise communication constraints as described in the previous section. In particular, we shall define which edges are activated at each step k and appropriately set the values for the α_k weights in the update rule (2) to achieve finite time convergence of all agents' estimates to the average (1).

We will first address the case when the number of agents in the ring is even, and then we will extend the result to rings with an odd number of agents.

A. Ring with an even number of agents

For a ring with an even number of agents ($m = 2n$), the coefficients of the linear combination can be chosen according to the following rule:

$$\alpha_k = \begin{cases} \frac{k}{k+1}, & 1 \le k < n \\ \frac{1}{2}, & k = n \\ 0, & k > n \end{cases} \qquad (3)$$

which due to (2) entails that the estimate is kept constant after $k = n$ steps and hence the iterative process can be terminated.

Algorithm 1 Algorithm for agent i – m even

1: $x_i(0) \leftarrow$ initial value for agent i
2: **for** $k = 1$ **to** n **do**
3: *% Select a neighbor to communicate with*
4: **if** $i + k$ is even **then**
5: $j \leftarrow \mathrm{Post}(i)$
6: **else**
7: $j \leftarrow \mathrm{Pre}(i)$
8: **end if**
9: *% Update the estimate using* (2) *and* (3)
10: $x_i(k) \leftarrow (1 - \alpha_k)x_i(k-1) + \alpha_k x_j(k-1)$
11: **end for**
12: **return** $x_i(n)$

The distributed overall procedure is obtained by making each agent run synchronously with the others Algorithm 1, which involves alternatively communicating with the predecessor and successor neighbor in the ring topology (according to the agent indexing shown in Figure 1), while updating its estimate according to (2) with the coefficients of the convex combination defined as in (3). Agents should agree on the neighbor to communicate with at the first iteration, i.e., which of the two edge groups of Figure 1 should be simultaneously activated. This choice is actually embedded in Algorithm 1 (see line 4) and it is simply based on the agent identification number i. Functions $\mathrm{Post}, \mathrm{Pre}$: $\{1, 2, \ldots, 2n\} \to \{1, 2, \ldots, 2n\}$ in Algorithm 1 serve the purpose of specifying the successor and predecessor neighbor, respectively, and are given by:

$$\mathrm{Post}(i) = \begin{cases} i + 1, & i = 1, \ldots, 2n - 1 \\ 1, & i = 2n \end{cases} \qquad (4)$$

and

$$\mathrm{Pre}(i) = \begin{cases} 2n, & i = 1 \\ i - 1, & i = 2, \ldots, 2n. \end{cases} \qquad (5)$$

According to Algorithm 1, at iteration $k = 1$ agent $i = 1$ communicates with its successor agent 2, and, hence, the blue colored edges in Figure 1 are activated.
The following theorem holds for the proposed distributed scheme.

Theorem 1 (Finite time consensus)**.** *Given a ring network with $m = 2n$ agents, if all of them apply Algorithm 1 synchronously, then, after n iterations, the estimate x_i computed by each agent i, $i = 1, 2, \ldots, 2n$, equals the average of their initial values* (1)*.*

Proof: Our aim is to show that by applying Algorithm 1 to a ring with $2n$ agents we get

$$x_i(n) = \frac{1}{2n} \sum_{j=1}^{2n} x_j(0), \qquad (6)$$

for every agent i, $i = 1, 2, \ldots, 2n$.

It suffices to show that (6) *holds for i such that $i + n$ is even (i.e., i even if n is even, odd if n is odd). This is due to the fact that Algorithm 1 agent $i + 1$ (for which $i + 1 + n$ is*

odd) communicates with agent i at step $k = n$ and its estimate satisfies

$$x_{i+1}(n) = (1-\alpha_n)x_{i+1}(n-1) + \alpha_n x_i(n-1),$$

which is equal to $x_i(n) = (1-\alpha_n)x_i(n-1) + \alpha_n x_{i+1}(n-1)$ given that $\alpha_n = \frac{1}{2}$. Then, all pairs of agents $(i, i+1)$ will converge to the average in n steps.

We shall then focus next on agent i such that $i+n$ is even, and prove that the following equation hold

$$x_i(n) = \frac{1}{2n}\sum_{j=i-s+1}^{i+s} x_j(n-s-1) + \frac{n-s}{2n}\Big[x_{i-s}(n-s-1) + x_{i+s+1}(n-s-1)\Big], \quad (7)$$

for all $s = 1, \ldots, n-1$. Note that, by substituting $s = n-1$ in (7)*, we get* (6)*, which concludes the proof of the theorem.*

The proof of equation (7) *is by induction, i.e., we show that it is true for $s = 1$ (step 1), assume that it holds for some s (induction hypothesis), and then show that this is also the case for $s+1$ (step 2).*

Step 1: *By Algorithm 1,*

$$x_i(n) = \frac{1}{2}x_i(n-1) + \frac{1}{2}x_{i+1}(n-1). \quad (8)$$

At step $k = n-1$, we have that $i+k = i+n-1$ is odd and $i+1+k = i+n$ is even, hence

$$x_i(n-1) = \frac{1}{n}x_i(n-2) + \frac{n-1}{n}x_{i-1}(n-2), \quad (9)$$

$$x_{i+1}(n-1) = \frac{1}{n}x_{i+1}(n-2) + \frac{n-1}{n}x_{i+2}(n-2), \quad (10)$$

and, by substituting (9) *and* (10) *into* (8)*, we obtain*

$$x_i(n) = \frac{1}{2n}\Big[x_i(n-2) + x_{i+1}(n-2)\Big] + \frac{n-1}{2n}\Big[x_{i-1}(n-2) + x_{i+2}(n-2)\Big].$$

which is (7) *with $s = 1$.*

Step 2: *For the sake of clarity we will treat the two terms in* (7) *separately: the summation first and then the other term.*

First term: *Let us start by considering the first two contributions in the summation, that is:*

$$x_{i-s+1}(n-s-1) + x_{i-s+2}(n-s-1). \quad (11)$$

By Algorithm 1 when $k = n-s-1$ (and, hence, $i-s+1+k = i+n-2s$ is even and $i-s+2+k = i+n-2s+1$ is odd), we get

$$x_{i-s+1}(n-s-1) = \frac{1}{n-s}x_{i-s+1}(n-s-2) + \frac{n-s-1}{n-s}x_{i-s+2}(n-s-2), \quad (12)$$

$$x_{i-s+2}(n-s-1) = \frac{1}{n-s}x_{i-s+2}(n-s-2) + \frac{n-s-1}{n-s}x_{i-s+1}(n-s-2). \quad (13)$$

Substituting (12) *and* (13) *into* (11) *we get that*

$$x_{i-s+1}(n-s-1) + x_{i-s+2}(n-s-1) = x_{i-s+1}(n-s-2) + x_{i-s+2}(n-s-2). \quad (14)$$

The time-invariance property for the pairwise summation in (14) *holds true for all other pairs in the summation of the first term in* (7)*, which always contains an even number of terms, i.e., $2s$. We can thus conclude that*

$$\sum_{j=i-s+1}^{i+s} x_j(n-s-1) = \sum_{j=i-s+1}^{i+s} x_j(n-s-2). \quad (15)$$

Second term: *Now consider the second term in* (7)*, recalled here for ease of reference:*

$$\frac{n-s}{2n}\Big[x_{i-s}(n-s-1) + x_{i+s+1}(n-s-1)\Big]. \quad (16)$$

By Algorithm 1 when $k = n-s-1$ (and, hence, $i-s+k = i+n-2s-1$ is odd, while $i-s+1+k = i+n-2s$ is even), we get

$$x_{i-s}(n-s-1) = \frac{1}{n-s}x_{i-s}(n-s-2) + \frac{n-s-1}{n-s}x_{i-s-1}(n-s-2), \quad (17)$$

$$x_{i+s+1}(n-s-1) = \frac{1}{n-s}x_{i+s+1}(n-s-2) + \frac{n-s-1}{n-s}x_{i+s+2}(n-s-2). \quad (18)$$

Substituting (17) *and* (18) *into* (16)*,* (16) *reduces to*

$$\frac{1}{2n}\Big[x_{i-s}(n-s-2) + x_{i+s+1}(n-s-2)\Big] + \frac{n-s-1}{2n}\Big[x_{i-s-1}(n-s-2) + x_{i+s+2}(n-s-2)\Big]. \quad (19)$$

Finally, substituting (15) *and* (19) *into* (7)*, , which we assume that holds due to our induction hypothesis, we have that*

$$x_i(n) = \frac{1}{2n}\sum_{j=i-s}^{i+s+1} x_j(n-s-2) + \frac{n-s-1}{2n}\Big[x_{i-s-1}(n-s-2) + x_{i+s+2}(n-s-2)\Big],$$

which is the same expression as (7) *with $s+1$ in place of s, thus concluding the proof by induction of* (7)*.*

■

Remark 1 (Speed of convergence)**.** Since the diameter of a ring network with $2n$ agents is n, we need at least n iterations to make sure that each agent receives information from all the other agents. This implies that the result of Theorem 1, which shows that we need exactly n iterations to converge to the average in ring networks with an even number of agents, is tight. To the best of our knowledge, this outperforms existing results in the literature.

Remark 2 (Efficiency)**.** Since at each iteration every agent communicates with a single agent only, the information

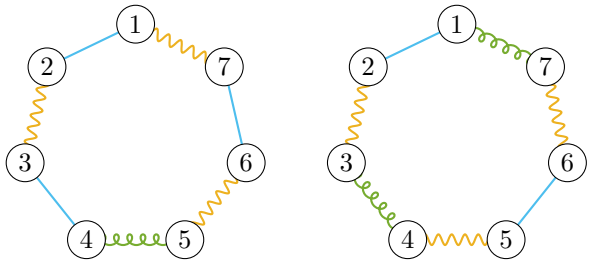

Fig. 2. Examples of communication protocols in a ring with an odd number of agents ($m = 7$): edges with the same colors are activated simultaneously. Three rounds are needed to activate all edges in both protocols.

transmitted in the proposed distributed algorithm is limited compared to alternative solutions in the literature with either fixed or time-varying topology without gossip constraints.

Remark 3 (Interpretation)**.** At time k, the update of agent i, $i = 1, \dots, m$, in (2) can be alternatively written as

$$x_i(k) = x_i(k-1) - \alpha_k(x_i(k-1) - x_j(k-1)). \qquad (20)$$

By inspection of (20), it can be observed that the estimate of the average that each agent maintains evolves as a discrete time integrator, where the quantity that gets integrated is the mismatch between the communicating agents' estimates, weighted by α_k. This can be thought of as the evolution of a closed-loop dynamical system, with the feedback gain fixed according to (3). The interpretation of Theorem 1, is that the dynamical system (20), reaches $\bar{x}$ at the n-th step, for all $i = 1, \dots, m$.

B. Ring with an odd number of agents

In the case of an odd number $m = 2n + 1$ of agents, we have to define a proper sequence of mode activations, thus choosing a coloring scheme for the ring. As stated earlier in Section II, for the odd case we need at least three colors for all edges to be activated in the minimum number of rounds while satisfying the gossip constraint. Whilst in the case of a ring with an even number of agents the coloring scheme is essentially unique (the solutions where two colors are swapped are indeed equivalent, see Figure 1), with three colors we can have multiple coloring configurations. To better clarify this observation, in Figure 2 we report two examples of a coloring scheme with three colors (cyan, yellow and green) for $m = 7$, which translates into two different multigossip sequences for the agents' communication. Interestingly, the solution that we propose in this paper does not depend on the adopted coloring scheme, which can be arbitrarily chosen but has to be agreed upon execution of the algorithm.

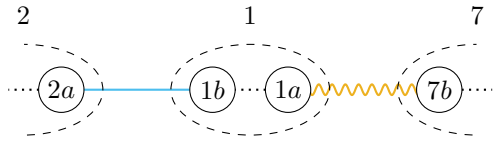

Fig. 3. Example of subagents and virtual communication links. Original agents are represented as dashed ellipsoids and virtual links as black dotted lines.

Algorithm 2 Algorithm for agent i – m odd

$x_{ia}(1) \leftarrow$ initial value for agent i
$x_{ib}(1) \leftarrow$ initial value for agent i
$k \leftarrow 2$
while $k \leq m$ **do**
  % Exchange information with neighbors
  $j \leftarrow \mathrm{Pre}(i)$
  $x_{ia}(k) \leftarrow (1 - \alpha_k)x_{ia}(k-1) + \alpha_k x_{jb}(k-1)$
  $j \leftarrow \mathrm{Post}(i)$
  $x_{ib}(k) \leftarrow (1 - \alpha_k)x_{ib}(k-1) + \alpha_k x_{ja}(k-1)$
  $k \leftarrow k + 1$
  % Update local estimates
  $x_{ia}(k) \leftarrow (1 - \alpha_k)x_{ia}(k-1) + \alpha_k x_{ib}(k-1)$
  $x_{ib}(k) \leftarrow (1 - \alpha_k)x_{ib}(k-1) + \alpha_k x_{ia}(k-1)$
  $k \leftarrow k + 1$
end while
return $x_{ia}(m)$

The solution for the case of an odd number of agents is based on the idea of making the number of agents even by considering each agent i as if it were composed by two subagents, say ia and ib, which can communicate with the predecessor and successor neighbor of agent i respectively, and are connected together by a *virtual communication link*. The communication link is virtual because it does not require any actual communication, but just a computation step for agent i. Figure 3 represents a pictorial view of this principle with reference to agent $i = 1$ of the ring network at the left panel of Figure 2. Original agents are represented as dashed ellipsoids and virtual links as black dotted lines.

Clearly the topology of this virtual network is still a ring, but it now has an even number $2m$ of agents. Moreover, if we set

$$x_{ia}(0) = x_{ib}(0) = x_i(0), \qquad (21)$$

then the average of the initial values of the $2m$ subagents equals the average of the initial values of the original m agents. By relabeling the subagents with numbers from 1 (for subagent 1a) to $2m$ (for subagent mb), one can recast this problem to the framework of Section III-A. Thus, by making each subagent run Algorithm 1 synchronously with the others, according to Theorem 1 we should obtain the average in a finite number of steps m, by alternating the edge activation as in Figure 1 with $2m$ in place of m, with the blue straight links being replaced by the black dotted lines playing the role of the virtual links, and the red wavy links being thought of as the actual communication links, which with reference to Figure 2, would be the cyan, the yellow and the green ones.

Following this principle, however, there is a remaining issue that needs to be addressed, that refers to the fact that simultaneously activating the cyan, yellow and green links would violate the pairwise communication constraint. This can be alleviated by activating them not all at the same time, but rather following a suitable communication scheme that is compliant with the gossip constraint. For example, we can first activate the cyan straight edges, then the yellow wavy ones,

and finally the green spring edges. This leads to an admissible solution but increases the number of communication rounds, that, unlike the even case, are no longer in correspondence to the algorithm iterations, but rather a multiple of 3. In particular, by Theorem 1 we know that if the agents in the ring are $2m$, then the average is achieved in m steps. Let $m = 2n + 1$. If the virtual communication links are activated first, then we have $n + 1$ iterations corresponding to virtual communication links and n iterations corresponding to the presence of an actual communication link. However, following the preceding discussion, the number of communication rounds corresponding to the presence of an actual link would be $3n$, i.e., it is a multiple of 3 of the corresponding number of iterations, to account for the gossip constraint.

Remark 4. Due to the fact that $\alpha_1 = 1/2$ (see (3)), if we start by activating the virtual links, then the first iteration of the algorithm ($k = 1$) is not needed since it computes the average of the initial estimates of the two subagents composing each agent which are set equal according to (21).

Algorithm 2 describes the steps performed by agent i. Note that each agent maintains two estimates that are updated based on the information received from its successor and predecessor neighbor and are then combined when the virtual edge is activated.

IV. Conclusions

In this paper we proposed an algorithm for finite time distributed averaging in the case of a ring network of agents, subject to a gossip constraint on communications. Interestingly, if the number of agents is even, then, consensus to the actual average is achieved in the minimum possible number of iterations, i.e., the diameter of the network, whereas the number of iterations needed in the case where the number of agents os od is higher, but still finite. Besides being of interest on its own, the proposed algorithm can also be embedded in distributed optimization schemes where computing the average is instrumental to solving the optimization problem (see e.g. [17]) and finite time convergence is then a requirement. These optimization schemes has been further developed and tailored to large scale systems in new application areas like energy system [18], which could benefit from a distributed finite time averaging algorithm.

References

[1] J. N. Tsitsiklis, "Problems in decentralized decision making and computation," Ph.D. dissertation, 1984.

[2] N. A. Lynch, *Distributed algorithms*. Morgan Kaufmann, 1996.

[3] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 988–1001, 2003.

[4] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Systems & Control Letters*, vol. 53, no. 1, pp. 65–78, 2004.

[5] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Gossip algorithms: Design, analysis and applications," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 3. IEEE, 2005, pp. 1653–1664.

[6] B. J. Liu, S. Mou, A. S. Morse, B. Anderson, and C. Yu, "Deterministic gossiping," *Proceedings of the IEEE*, vol. 99, no. 9, pp. 1505–1524, 2011.

[7] A. Olshevsky and J. N. Tsitsiklis, "Convergence speed in distributed consensus and averaging," *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 33–55, 2009.

[8] S. Sundaram and C. N. Hadjicostis, "Finite-time distributed consensus in graphs with time-invariant topologies," in *American Control Conference, 2007. ACC'07*. IEEE, 2007, pp. 711–716.

[9] C.-K. Ko and X. Gao, "On matrix factorization and finite-time average-consensus," in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*. IEEE, 2009, pp. 5798–5803.

[10] L. Georgopoulos, "Definitive consensus for distributed data inference," Ph.D. dissertation, ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE, 2011.

[11] A. Y. Kibangou, "Finite-time average consensus based protocol for distributed estimation over awgn channels," in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*. IEEE, 2011, pp. 5595–5600.

[12] J. M. Hendrickx, R. M. Jungers, A. Olshevsky, and G. Vankeerberghen, "Graph diameter, eigenvalues, and minimum-time consensus," *Automatica*, vol. 50, no. 2, pp. 635–640, 2014.

[13] S. Mou and A. Morse, "Finite-time distributed averaging," in *American Control Conference (ACC), 2014*. IEEE, 2014, pp. 5260–5263.

[14] J. M. Hendrickx, G. Shi, and K. H. Johansson, "Finite-time consensus using stochastic matrices with positive diagonals," *IEEE Transactions on Automatic Control*, vol. 60, no. 4, pp. 1070–1073, 2015.

[15] G. Oliva, R. Setola, and C. N. Hadjicostis, "Distributed finite-time average-consensus with limited computational and storage capability," *IEEE Transactions on Control of Network Systems*, no. 99, 2016.

[16] G. Shi, B. Li, M. Johansson, and K. H. Johansson, "Finite-time convergent gossiping," *arXiv preprint arXiv:1206.0992*, 2015.

[17] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: numerical methods*. Prentice-Hall, Inc., 1989.

[18] A. D. Dominguez-Garcia and C. N. Hadjicostis, "Distributed algorithms for control of demand response and distributed energy resources," in *2011 50th IEEE Conference on Decision and Control and European Control Conference*. IEEE, 2011, pp. 27–32.